Contribution ID: **9** Type: **not specified**

# Shared memory and message passing revisited in the many-core era

*Wednesday, 2 March 2016 14:00 (1 hour)*

In the 70s, Edsgar Dijkstra, Per Brinch Hansen and C.A.R Hoare introduced the fundamental concepts for concurrent computing. It was clear that concrete communication mechanisms were required in order to achieve effective concurrency.

Whether you're developing a multithreaded program running on a single node, or a distributed system spanning over hundreds of thousands cores, the choice of the communication mechanism for your system must be done intelligently because of the implicit programmability, performance and scalability trade-offs. With the emergence of many-core computing architectures many assumptions may not be true anymore.

In this talk we will try to provide insight on the characteristics of these communication models by providing basic theoretical background and then focus on concrete practical examples based on indicative use case scenarios. The case studies of this presentation cover popular programming models, operating systems and concurrency frameworks in the context of many-core processors.

**Presenter:** SANTOGIDIS, Aram (CERN)